LHC Software and Computing

Saul Gonzalez
DOE

HEPAP Meeting
Washington, DC
April 18-19, 2004

Outline

- LHC Research Program
  - Software and Computing
- Scope of Software and Computing
- US ATLAS & US CMS Computing Projects
  - Computing Model
  - Facilities Model
  - Management and Oversight
- Recent progress and outlook
  - Grids
  - Grid2003
  - Interactions with LCG
- Conclusions

US LHC Research Program

- NSF—DOE Partnership
  - Joint funding
  - Oversight through Joint Oversight Group (JOG)
- Major components:
  - Experiments
    - Maintenance and Operations (M&O)
    - Software and Computing (S&C)
    - Detector Upgrade R&D
  - Accelerator R&D
- Explicit S&C line in Research Program
  - Recognition that Software and Computing are key components to program success
  - Experiments can optimize split between M&O and S&C

US LHC Research Program / 2

- Research Program activities are ramping up sharply
  - Entering phase of pre-operations and detector exploitation

EXAMPLE: DOE Funding Profile

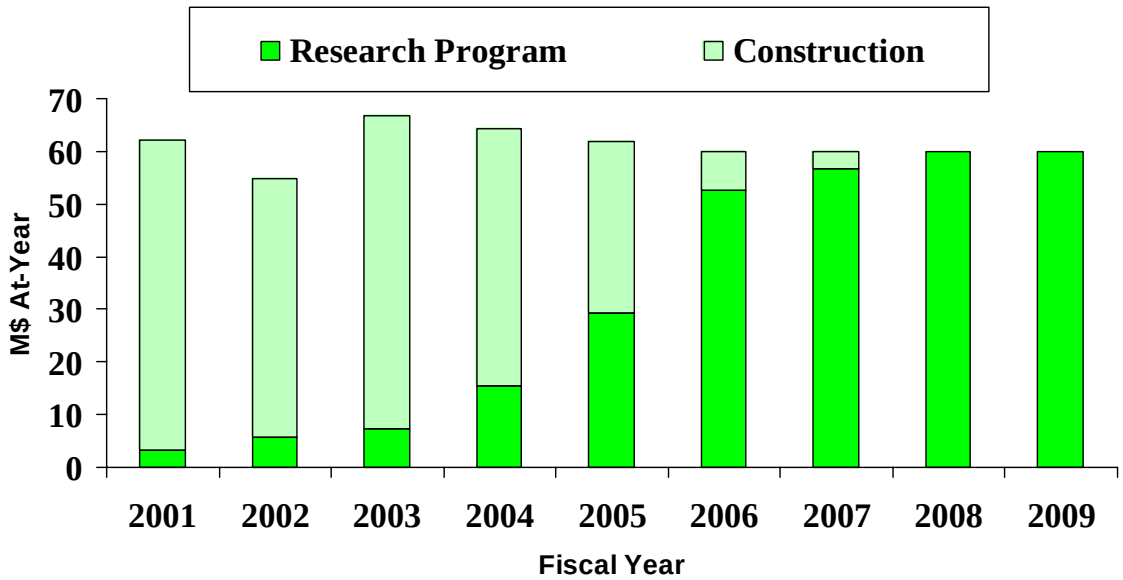

U.S. ATLAS and U.S. CMS Computing Projects

- Both experiments have defined Software & Computing Projects
- Recognition that S&C are critical for physics readiness
- Projections around turn-on are ~ $15M / experiment / year:
  - Roughly 2/3 personnel
  - Rest is hardware at Tier 1 and Tier 2 regional centers
- Detailed resource loaded schedules, milestones, etc.
- Project managers decide on priorities and allocate funds accordingly (as in construction project)
- This flexibility has allowed early hiring of significant number of software engineers, helping to place US in leadership position.

S&C Program Management

- U.S. Collaboration management teams are identified and vetted by the collaborations with concurrence of the agencies (JOG)
- Management exercises authority according to a Management Plan, which defines roles and responsibilities for individuals
  - ATLAS and CMS tailor their MP to their needs
- Recent changes in Management Plans reflect the shifting of activities from fabrication to research
- Oversight:
  - Frequent inter-agency and agency-management phone meetings
  - Quarterly progress reporting on milestones – used by agencies to monitor progress
  - External peer reviews

LHC S&C Reviews

- External peer reviews
  - Annual reviews with follow-up mini-review during the summer
  - Covering management, planning, resources, manpower, review of milestones and scope
- Provides useful input to the Collaborations
- Provides independent assessment of progress to the Agencies → Reviews are very positive oversight tool.
- Last review:
  - **Jan 13-16 at FNAL**
  - **Panel consisted of 9 outside reviewers (from US, Europe, and from labs and universities)**

More on last review later

Elements of LHC Software & Computing

- Different applications at different stages of the experiment:
  - Test beam
  - Data Challenges, validation of computing model
  - Cosmic runs, commissioning, testing
  - ➢ **Physics**
- Scope of U.S. LHC Software and Computing**:**
  - Sub-system reconstruction – natural involvement deriving from US detector responsibilities
  - Core software – framework with interfaces to services, data, algorithms
  - Services – geometry, calibration, alignment
  - General reconstruction and analysis algorithms
  - Grid-enabling software
- Important "hardware" aspect of S&C: Facilities

➢ Wide range of US involvement both in CMS and ATLAS

LHC S&C: Enabling University Science

- The success of the US participation in the LHC will ultimately depend on how effective US researchers are in extracting the physics
- The LHC computing model for the LHC experiments must therefore support location-independent access to LHC data and computing resources
  - It must bring LHC physics to the Universities, where a large fraction of the analyses are carried out
- The US must play a leading role in defining and developing this model

US Participation in LHC Computing

- U.S. participants have moved aggressively to take early leadership roles within the collaborations
- Leveraging expertise in the National Labs and Universities
- Strategy facilitates early entry into the Physics

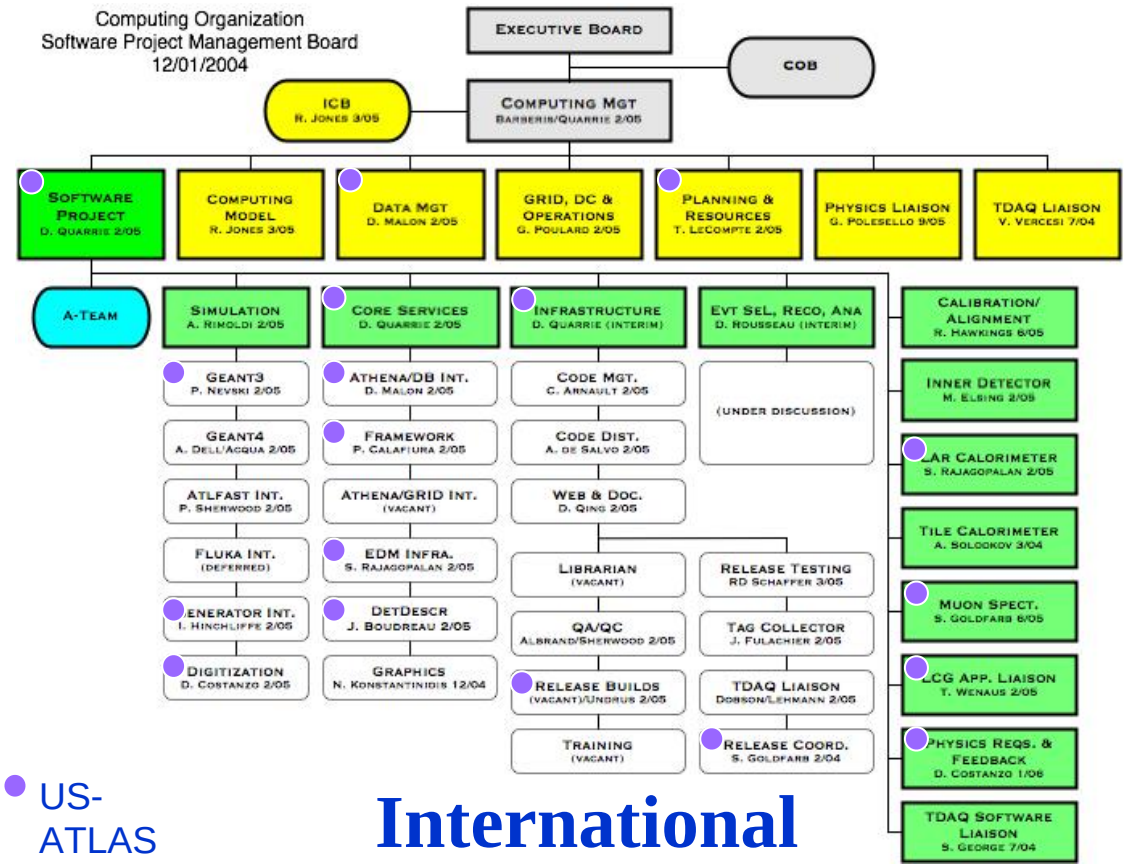

US Participation in LHC Computing / 2

International CMS

CCS = Computing and Core Software
PRS = Physics Reconstruction & Selection
TRIDAS = Trigger and DAQ Software

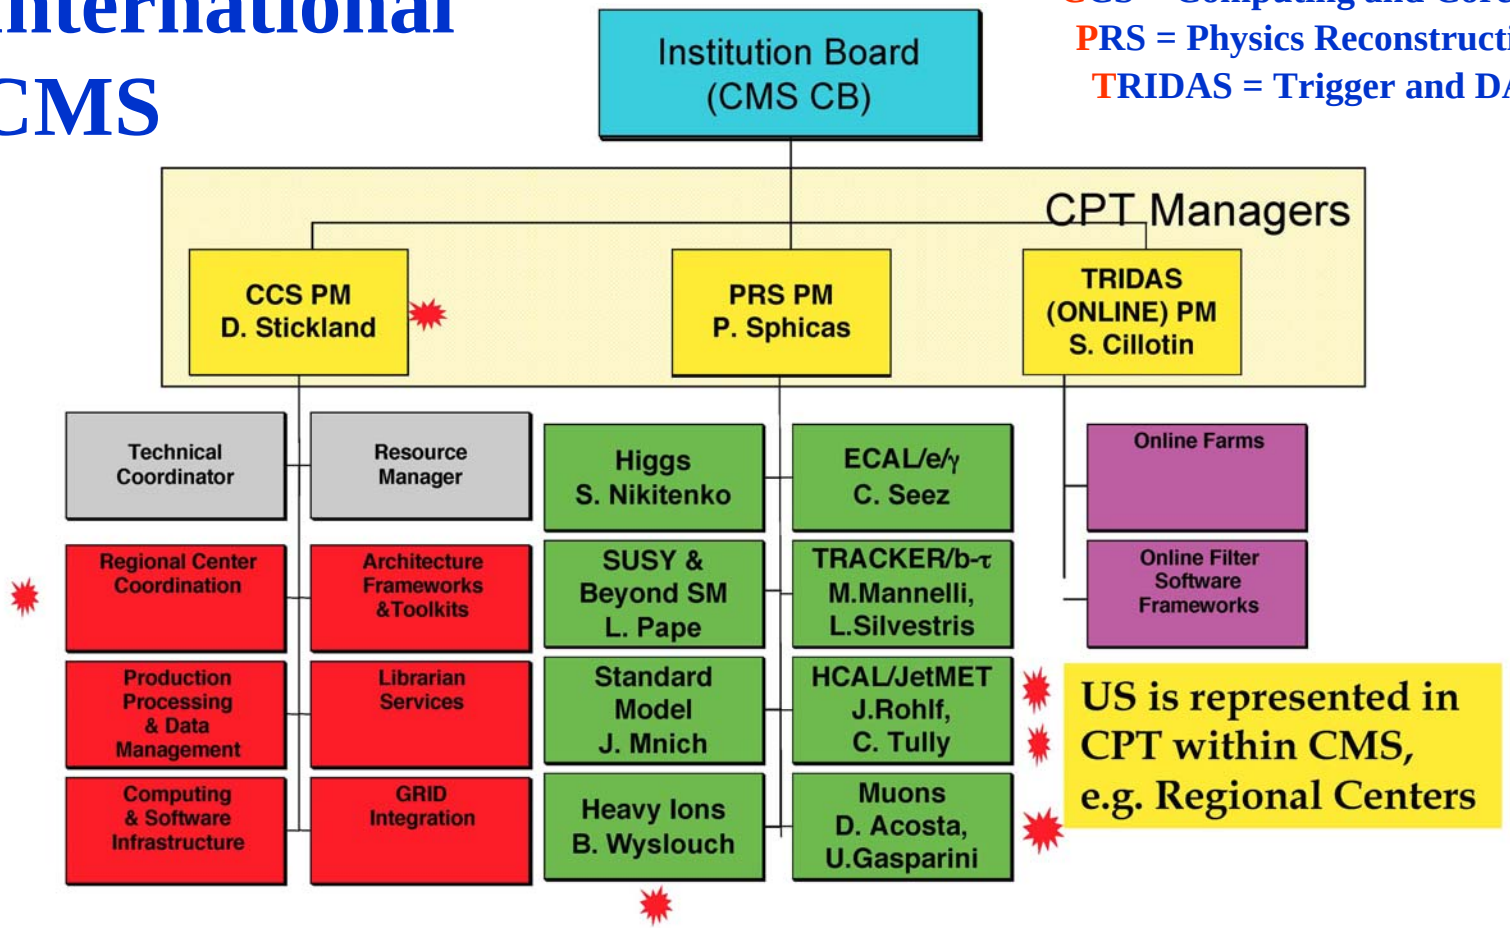

Software Professionals in LHC S&C

- Another metric: U.S. software professionals in CMS S&C
- Similar picture in ATLAS

- **Software professionals are essential** to address the complexity of modern computing.
  - For example, the dual use of offline software in the high level triggers requires a new level of discipline in software engineering.
- **No Moore's law for personnel**
  - While CPU and storage costs go down in time, personnel costs go up
  - S&C Program costs will be dominated by personnel, not facilities

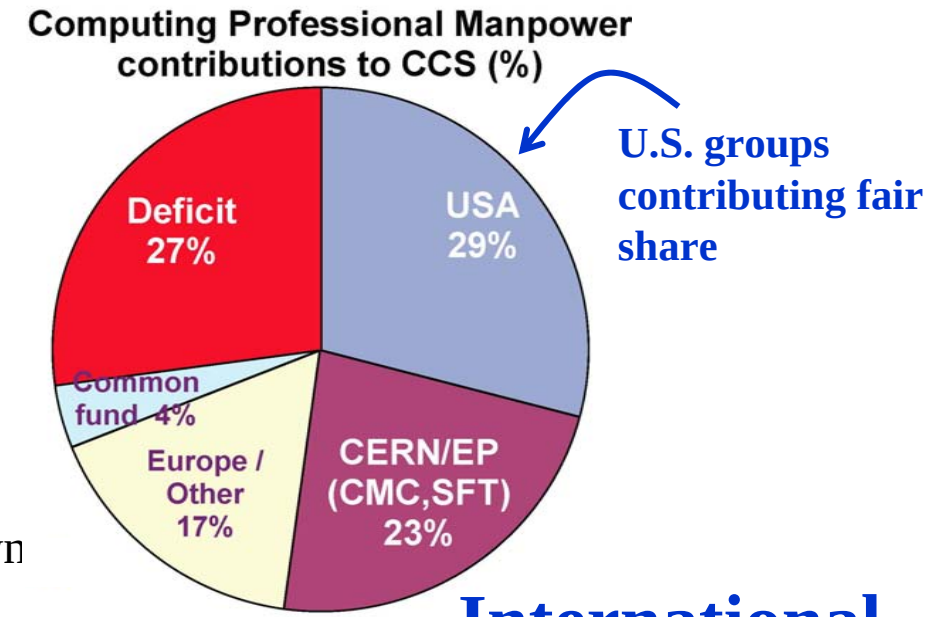

International CMS

Computing Model

- The LHC Computing model must support
  - Production (reconstruction, simulation)
  - Data Analysis
  - Code development and testing
- Only sensible way to handle large flow of data and need for CPU cycles at the LHC is to have a distributed computing model
  - Facilitates collaboration, by providing direct participation of remote sites
  - With the right interfaces, opens the door to opportunistic computing
    - Use idle resources not dedicated to the LHC
- Distributed model from the start (distributed resources and coherent global access to data)
- Success of this model is particularly important for U.S. physicists, who will benefit from transparent access to data and services
  - Equal access to data + Leadership → Good return on investment

Distributed Facilities Model

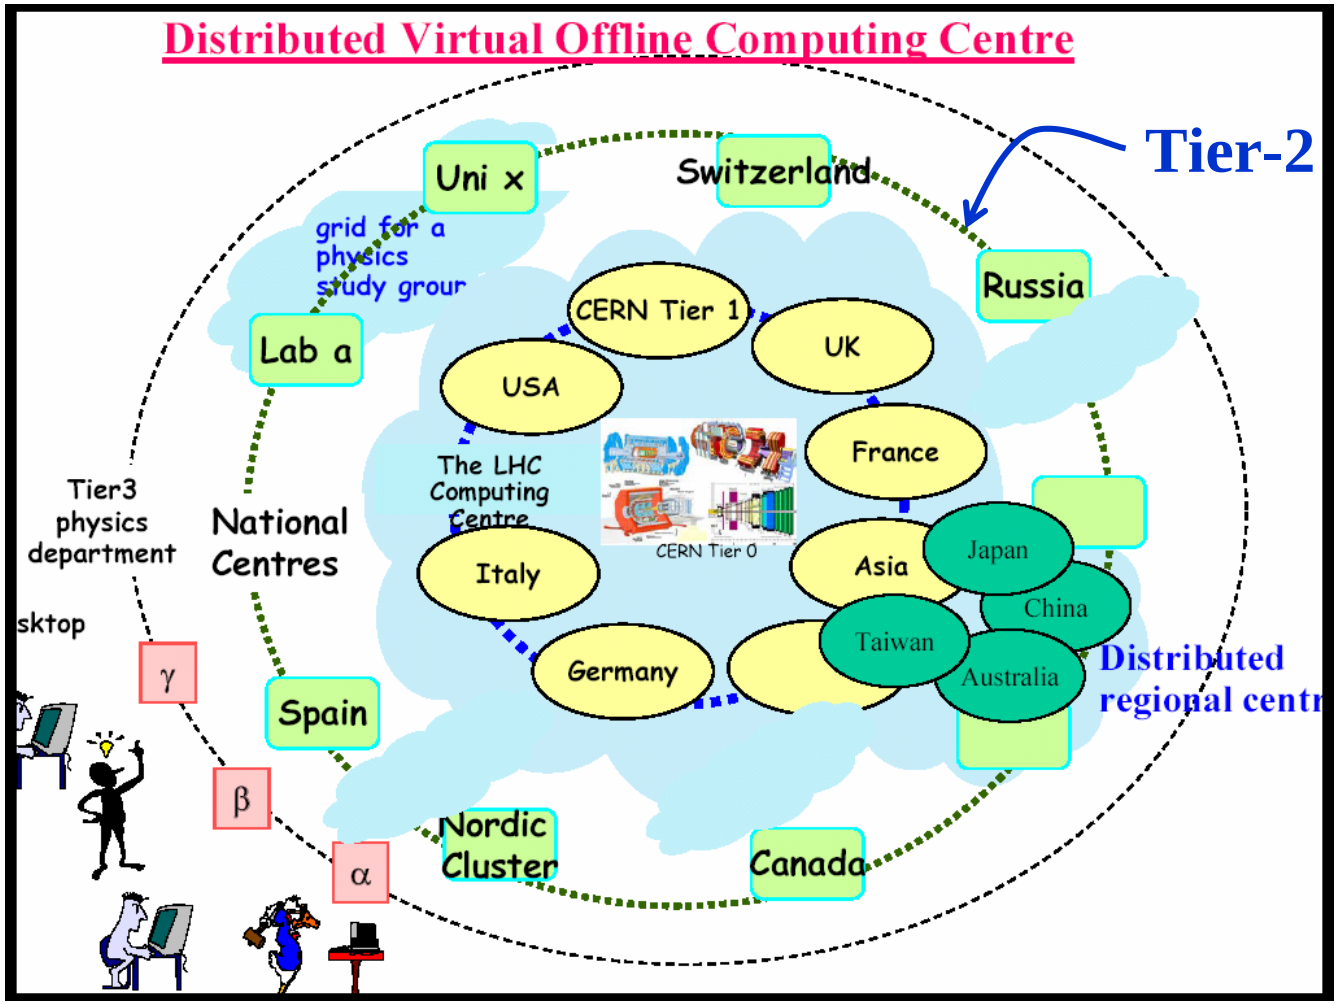

Distributed Facilities Model / 2

- TIER-0:
  - CERN computing center collecting raw data (interface to online)
  - Distributes data to regional centers; ~20% of all CPU resources
- TIER-1:
  - Provides grid-enabled CPU cycles and persistent/transient storage
  - Will commit to provide: adequate bandwidth, QoS, 24/7 services and support, long-term access to data over lifetime of LHC
  - Responsible for a consortium of Tier-2 centers under it
  - Expect ~7 centers worldwide with ~40% of all grid resources; in US:
    - BNL (ATLAS)
    - FNAL (CMS)
- TIER-2:
  - Provides grid-enabled CPU cycles and storage
  - 24/7 service
  - Expect ~40% of total grid resources to be at Tier-2
  - Enabling universities; backbone of distributed model
- TIER-3:
  - Universities, users

Details of this model are presently being defined/discussed

Grids

- LHC characterized by:
  - Large computing and data movement requirements
  - Large, world-wide collaborations
- Bring the data and CPU cycles to the researcher → Grid
- Grid paradigm is a lot more than distributed computing
  - A collection of global computing and storage resources operating under a global operating system

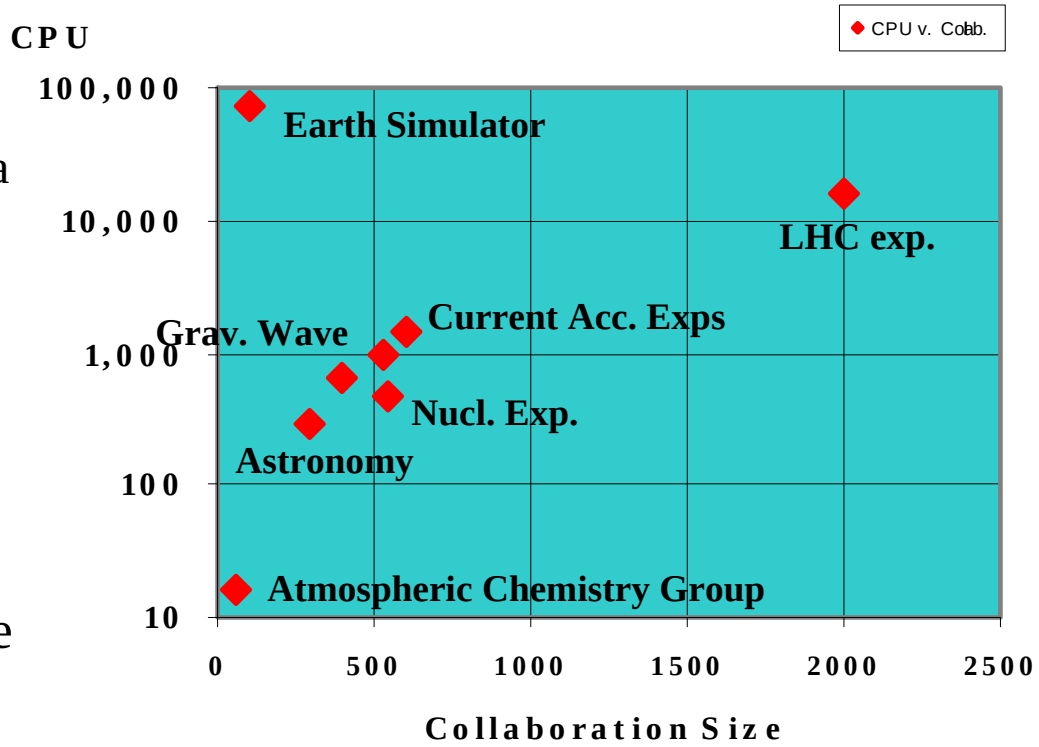

Key Concept: *Interoperability* allows different Grids to aggregate resources and services and to behave as a single entity.

Groundwork for a HEP Grid

- HEP-driven Grid research projects have made important contributions to the development of Grid software
  - PPDG (DOE)
  - GriPhyN (NSF)
  - iVDGL (NSF)
- Complementary projects; now coordinating convergence of developments under Trillium Project
  - Focused on integration of applications on the Grid
  - Provides single "Grid voice" in dealing with non-U.S. Grid projects
  - Real Grid deployment, driven by HEP applications
- Biggest success so far → Grid2003 **(See P.Avery's report at last HEPAP meeting)**
  - Substantial Grid deployment for real applications
    - Multi-application: LIGO, SDSS, CDF/D0, CMS/ATLAS
  - Beginnings of a Persistent Grid
- Long term support issues

Grid2003 Demonstrator

Grid2003 Project follow-on of US Atlas and US CMS Grid testbeds

- Demonstration for SC2003 and U.S. funding agencies: performance demonstrator for functional multi-VO Grid
- Collaboration of US LHC and Grid projects, labs and universities Including both U.S. Tier-1 and all U.S. Tier-2 centers

VO = Virtual Org.

Grid2003 approach

- experiment projects/VOs (US CMS, US Atlas and others) bring their grid-ified applications into multi-VO Grid3 environment
- Grid2003 team works with sites to provide basic Grid services:
  - processing and data transfer, software packaging/deployment, monitoring, information providers, VO/authentication management, basic policies
  - simple/non-intrusive installation based on VDT and EDG middleware
  - iVDGL iGOC cross-VO operations support, including trouble tickets

28 sites, 2800 CPUs, running fairly stable since SC2003 (Nov 2003)

- e.g., 13M CMS full detector simulation events produced on Grid3 -- and counting
- represents about 100 processor years of computing

CMS Data Challenge Underway

Pre-challenge production of 70M fully simulated events (20M w/ Geant-4)

- ➔ massive production of event samples during 2003/04
  - Large samples of simulated CMS events for DC04 also feed into physics studies for Physics Technical Design Report
- ➔ through U.S. Grid3, large CPU resources are made available to CMS
  - U.S. LHC Tier-1 and Tier-2 centers, Universities and other centers

DC04 full-chain demonstration of CMS data reconstruction (March-April 04)

- ➔ sustained data reconstruction at 25Hz at CERN Tier-0 center
  - corresponds to 25% of target conditions at LHC startup in 2007
- ➔ reconstructed data get transferred to Tier-1 centers
  - Fermilab for U.S. CMS, and UK, Italy, Spain, Germany, France Tier-1 centers
- ➔ management of the distributed CMS data sets on the LHC Grid
- ➔ running of fake analyses on DC04 data at Tier-1 and Tier-2 centers
- ➔ monitoring/archiving performance criteria for debugging and post mortem

Get wealth of information as input in Computing Technical Design Report

- ➔ Major milestone to get software and computing ready for CMS

ATLAS Data Challenge 2

- ATLAS DC2 starting in 2 weeks
  - Will use ~1000 CPU's continuously for 2 months
  - Will opportunistically use more resources on Grid3+
- Based on PreDc2 exercise of last Nov., 2003.
  - PreDC2 was the basis for our SuperComputing 2003 demo
- Grid3 has been a very successful operation involving many disciplines/collaborators.
  - The evolved version (grid3+) will be the basis for ATLAS DC2

US ATLAS Datasets on Grid3

- Grid3 resources used
  - 16 sites, ~1500 CPUs exercised; peak ~400 jobs over three week period
- Higgs → 4 lepton sample
  - Simulation and Reconstruction
  - 2000 jobs (X 6 subjobs); 100~200 events per job (~ **200K events**)
  - 500 GB output data files
- Top sample
  - Reproduce DC1 dataset: simulation and reconstruction steps
  - 1200 jobs (x 6 subjobs); 100 events per job (**120K sample**)
  - 480 GB input data files
  - Data used by PhD student at U. Geneva

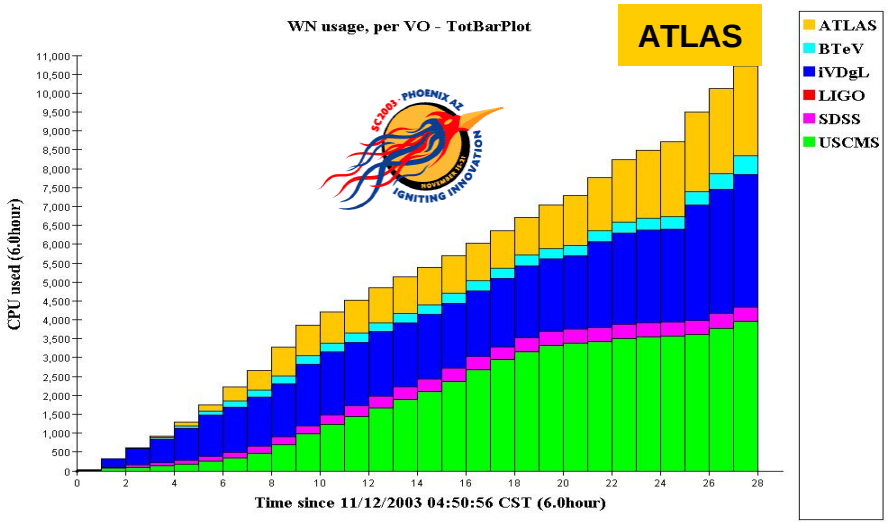

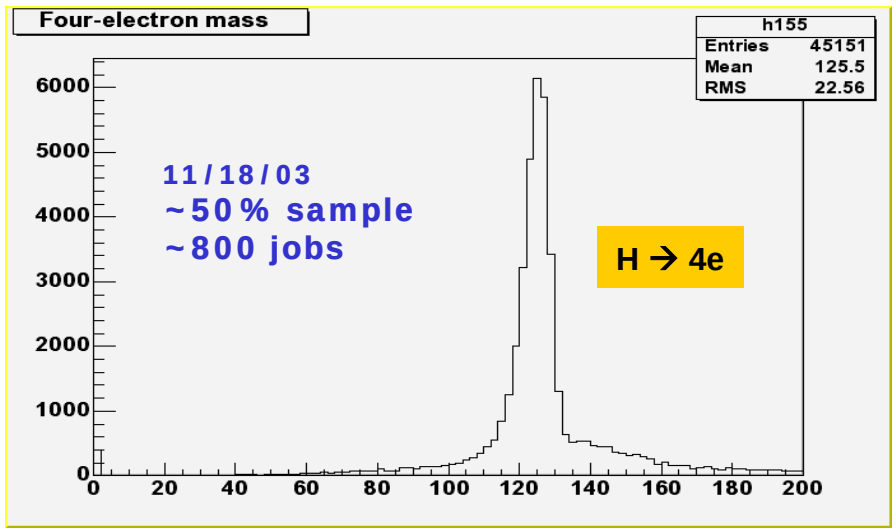

Toward the US **Open Science Grid**

Building partnerships on US Grid infrastructure for LHC and other sciences

- ➔ Federate currently disjoint grid resources into a single managed grid
- ➔ LHC application driving this effort, Grid3 is a great initial step
- ➔ Federate US resources with the LCG, the EGEE and other national and international Grids

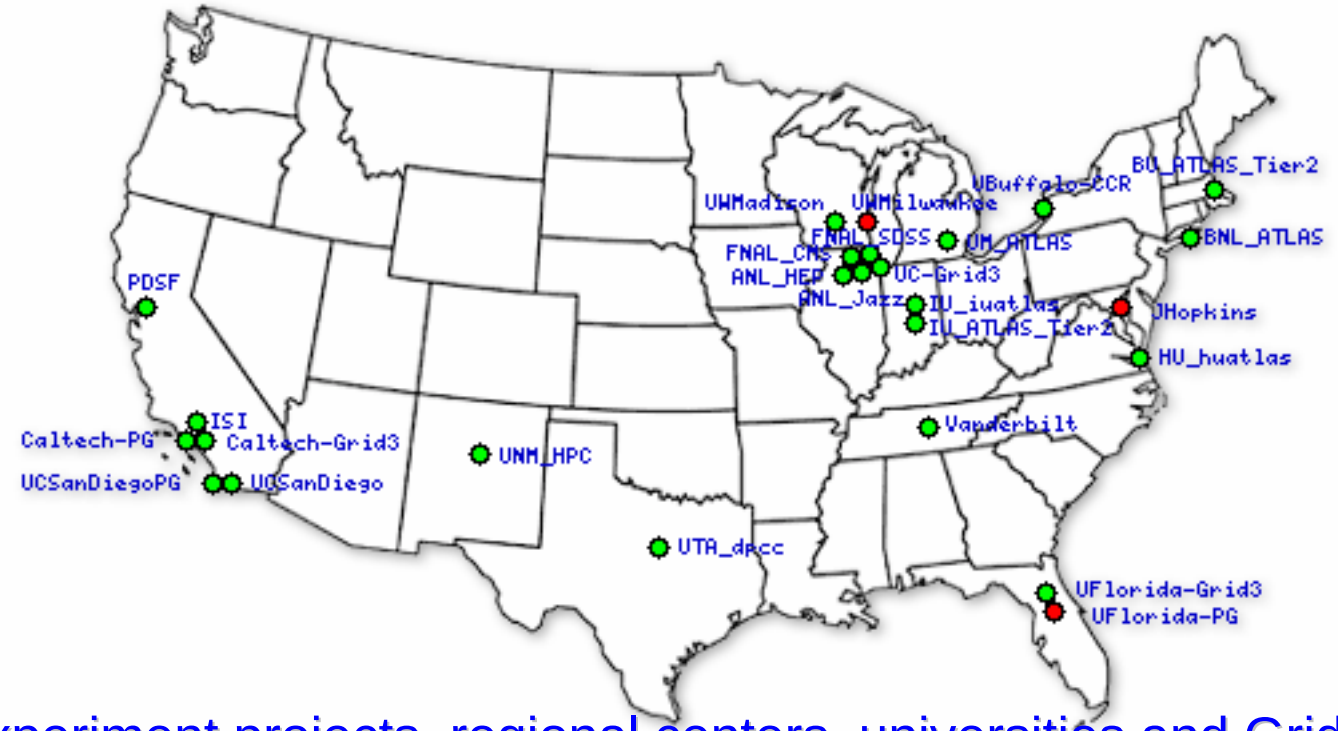

US LHC experiment projects, regional centers, universities and Grid projects formulated a roadmap towards the “Open Science Grid”

LHC Computing Grid (LCG)

- CERN-based LCG project is to prepare the grid infrastructure for LHC experiments
  - Many intellectual contributions from other parties – example: VDT middleware contributions (US)
- Project has two phases
  - Phase I (2002—2005):
    - Development of common LHC application support and services leading to the development of an LHC computing model.
    - Prototyping and validation of model via Data Challenges in the experiments.
    - Phase I culminates with the LCG TDR.
  - Phase II: (2006-2008): Gradual deployment of LCG
- LCG now driving EGEE, a European initiative on “e-science” grids
- LHC Computing is a world-wide distributed enterprise
  - Contributions from many regions, managed collaboratively
  - A new way of managing resources

LCG Organization

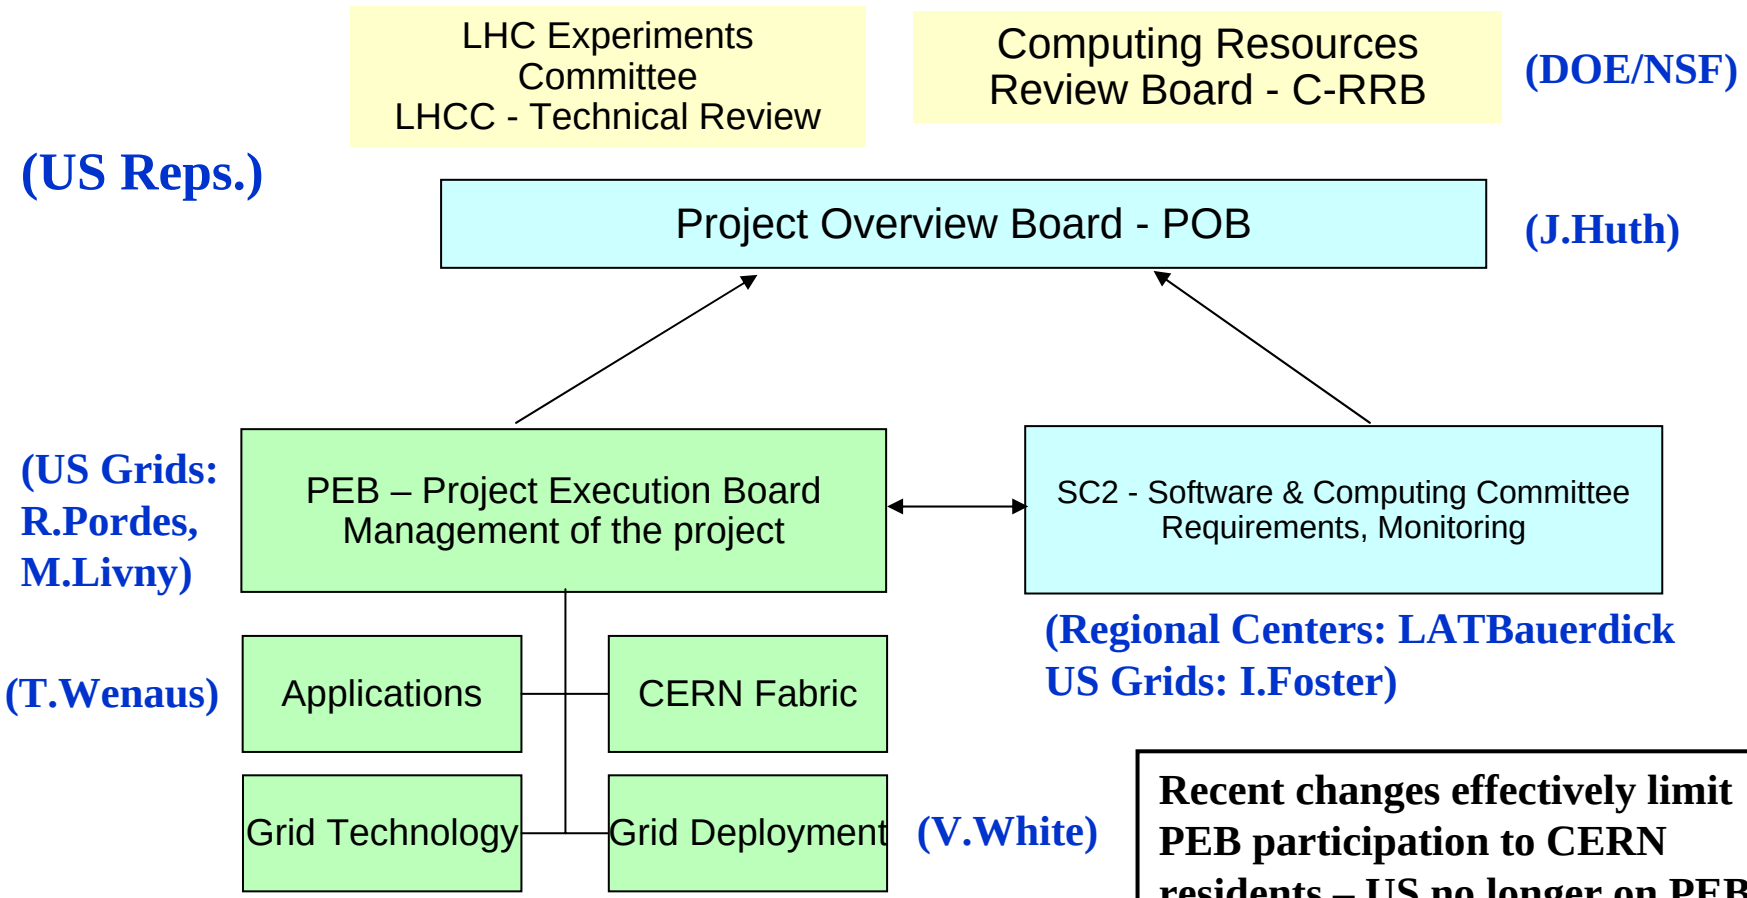

Recent changes effectively limit PEB participation to CERN residents – US no longer on PEB

Interoperation of US Grids with the LCG

US Atlas and US CMS working on interoperability of LCG and US Grid

- ➔ First steps already achieved
  - On storage service, middleware, VO management and application level
- ➔ Atlas DC2 application running across LCG, NorduGrid, US Grid3
- ➔ CMS DC04 data transfers and management of dataset replicas between storage services on LCG and US Grid3 sites

Next step: US Tier-1 centers to federate US resources with LCG service

- ➔ Realistic near term goals:
  - Fermilab Grid installation available to LCG resource broker through existing LCG-2 installation at Fermilab Tier-1
  - Reconciling LCG and US Grid VO management (VOMS)
- ➔ Next steps this year
  - Managed storage across Grids
  - Include access to US Tier-2 centers and other US Grid sites from LCG

Emerging Distributed Analysis (ARDA) approach to middleware and end-to-end systems will help in facilitating this

Federating Worldwide Resources for the LHC

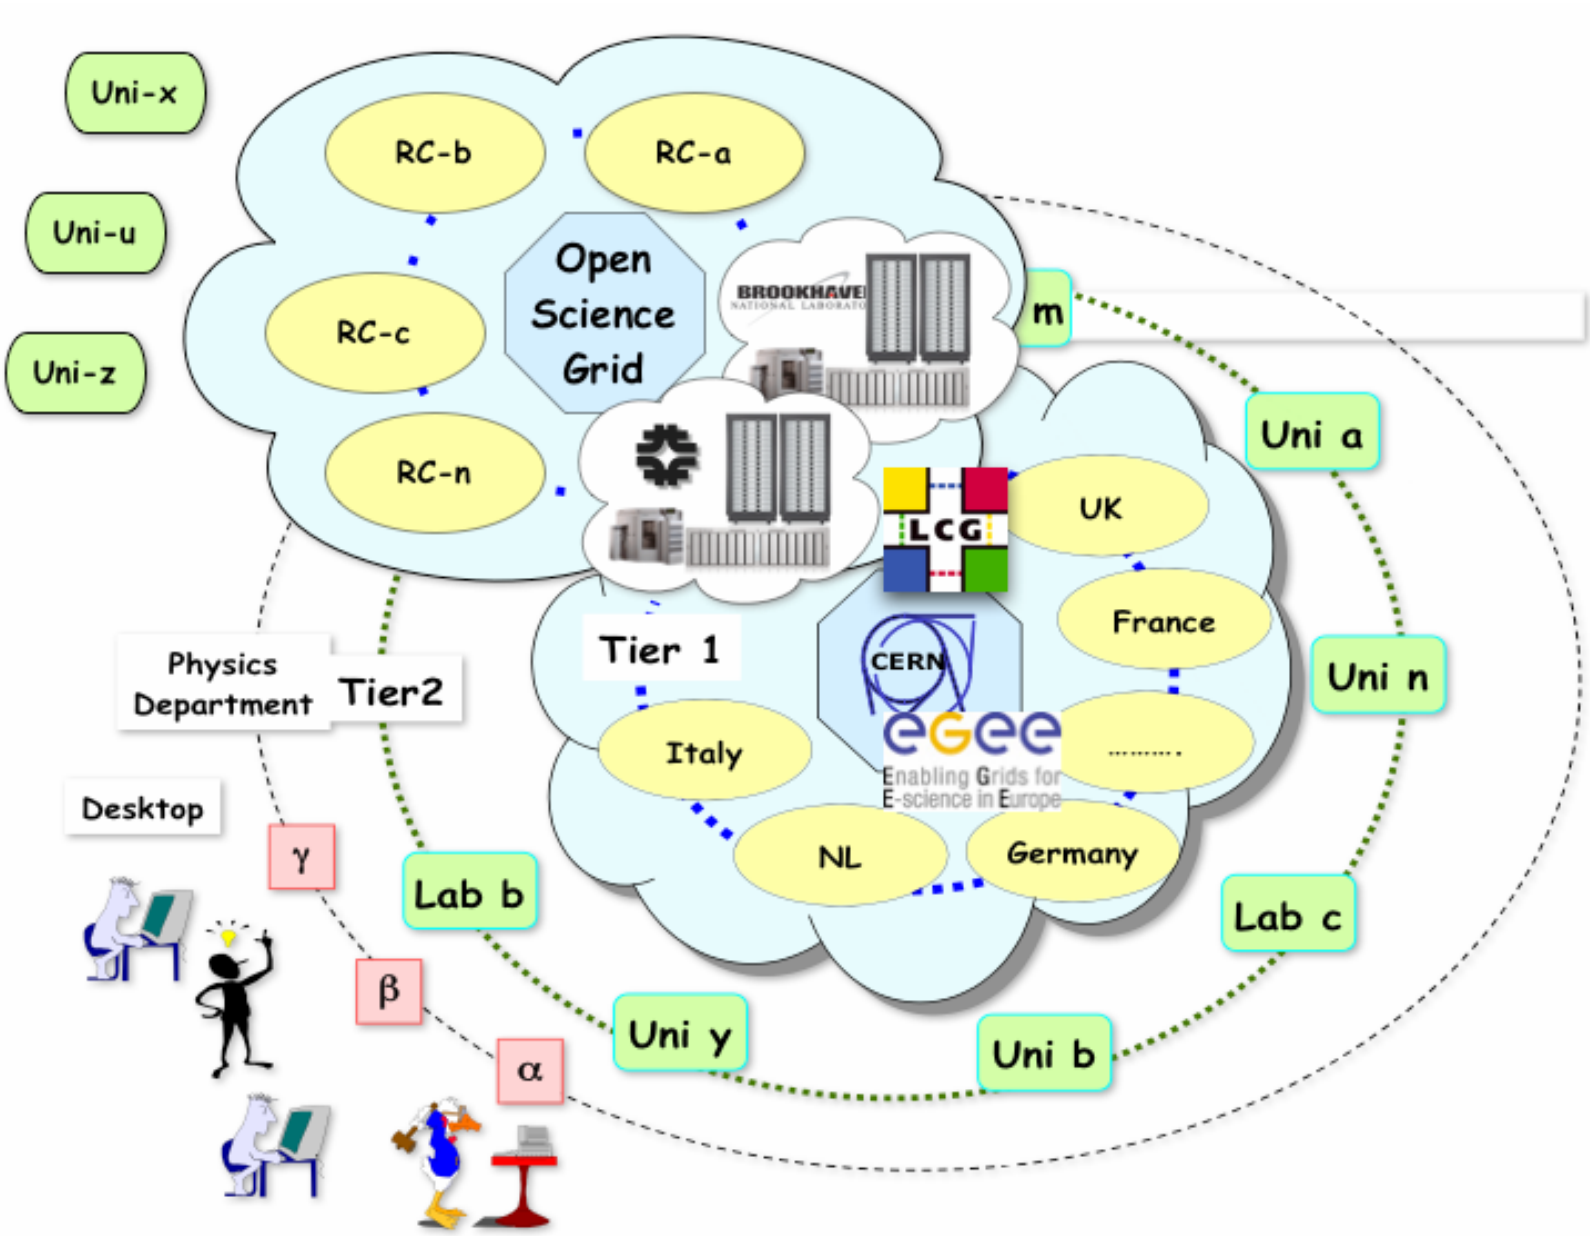

Challenges of Peering with the LCG

- Challenge for the LCG project to become fully aware of requirements of computing centers outside CERN and developing the appropriate management mechanisms for such a distributed enterprise
  - Until recently, LCG model was “central control of all LCG resources”
    - Makes it difficult to exploit resources in multi-experiment and multi-disciplinary computing centers
  - US strategy is to share resources among multiple experiments, multiple communities, as well as opportunistic resources (e.g., utilizing idle resources not dedicated to LHC)
    - This was the success of Grid2003
- Issue addressed at a meeting on April 7 at BNL between US-ATLAS, US-CMS, and LCG management
  - Agreement to collaborate in developing a roadmap for LHC computing that takes these and other technical challenges into account
- Interoperability is work in progress
  - Optimal solution is probably somewhere in between “proliferation of grids” and a monolithic grid → a few flavors of grid, federated
  - Prompt convergence on this issue is essential

Commitments to LCG

- Once the LCG project enters in production mode (Phase II), formal service level commitments will need to be established between the central facility (Tier-0 at CERN) and the major regional centers (Tier-1 and Tier-2)
  - Agreement between all entities that contribute services to the grid
- To that end, a task force was established to draft an LCG MoU between CERN and the participating funding agencies. Among other things, the MoU will cover:
  - Definition of Tier-0/Tier-1/Tier-2 relationship
  - Expected service levels
  - Organization
- MoU Status: Early drafts are being discussed by the task force and in the US LHC S&C community.

Other US LHC S&C Activities

- U.S. LHC Edge Computing Project
  - An installation of CPUs and storage servers at CERN that will facilitate high throughput data transfers between the Tier-0 and the Tier-1 facilities in the US.
  - The facility will also be made available to the LCG for tests of the Tier-0 computing model.
- US LHC contributions to LCG/EGEE middleware
  - U.S. is providing the technical leadership of the LCG/EGEE ARDA middleware project, which will build the first prototype for grid-enabled distributed analysis. This will insure coherence between the US and EGEE efforts.
- Contributions to EGEE Grid software deployment (NSF)
  - Effort to facilitate interoperability between US and European software. Emerging EGEE middleware will be factorized and packaged into future VDT releases, maintaining compatibility in the middleware layer.
- US contributions to LCG applications area
- Many, many contributions to CMS and ATLAS core, sub-system, and reconstruction software

US LHC Software & Computing External Review

- **Last Review dominated by concerns about potential funding shortfalls**
- **Projects are well managed; US-CMS and US-ATLAS are important contributors to international efforts**
- **Major issues:**
  - Budget cuts can no longer be covered by LHC schedule stretch out; any additional cuts will curtail US ability to do its share in upcoming data challenges, will compromise ability of US university physicists to do data analysis, and could force US to renege on international agreements
  - GRID2003 was tremendous success; US should continue to pursue strategy of grid interoperability to allow efficient exploitation of computing resources not "owned" by LHC and of multi-experiment and multi-science computing centers, and should build on success by development of plans for persistent grid infrastructure as proposed in the Open Science Grid
  - Base program erosion is beginning to have serious effects, especially at universities, and should be curtailed as much as possible

S&C Evolution to 2007

- Milestones
  - CMS DC04 (in progress)
  - ATLAS DC2 (May/2004)
  - LCG-2 Deployment
  - ATLAS Computing TDR (mid 2005)
  - CMS Computing TDR (December 2004)
  - LCG TDR (July 2005)
  - CMS DC05
  - CMS Physics TDR (End 2005)
  - ATLAS DC3 (End 2005)
  - ATLAS PRR (mid 2006)
  - ATLAS DC4
  - LCG production deployment (start in 2006)
  - Commissioning run (Fall 2006)

Conclusions

- An early and strong U.S. involvement in LHC S&C provides a strategic advantage that offsets some of the drawbacks of remote participation.
  - U.S. participation at all levels in both experiments and in the LCG
- The Software & Computing component of the Research Program will give U.S. physicists the tools and support necessary to obtain a commensurate return on the U.S. investment
  - Will enable Universities to have full access to LHC physics
  - Can't drop the ball now – must work to ensure adequate support
- Interoperable Grids is the right computing model for the LHC experiments
  - The U.S. is a leader in this area
- The LHC S&C is bringing together a large number of Universities and National Labs
  - Forging interagency, interdisciplinary, and international partnerships
  - Spurring new initiatives for a truly global science

Additional Slides

CMS running worldwide “Data Challenge”

Pre-challenge: massive Grid-based production of CMS simulated events at Universities, Tier-2 and Tier-1 centers across the U.S. Grid

DC04 Data Challenge: run a full chain of reconstruction, data distribution and analysis for a sustained period at 25% of LHC rate

GRID 2003

- US ATLAS PreDC2 exercise:
  - Development of ATLAS tools for DC2
  - Collaborative work on Grid2003 project
  - Gain experience with the LCG grid

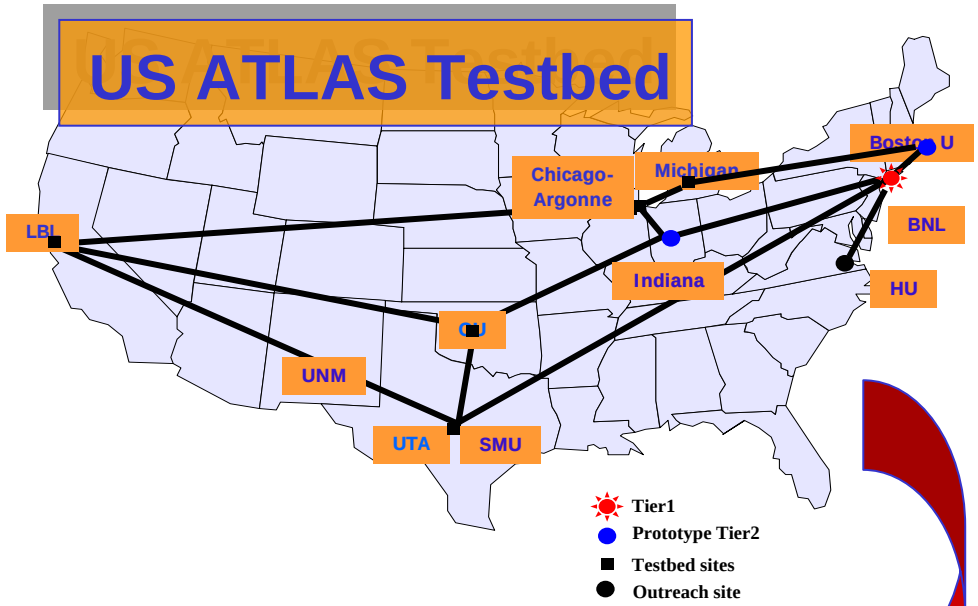

- Puts us in good shape for DC2

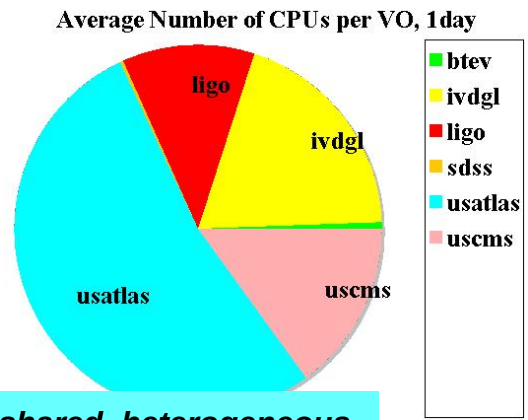

US ATLAS shared, heterogeneous resources contributed to Grid2003